First Annual Joint ISPA/SCEA International Conference

Toronto, Ontario, Canada
June, 1998

The JPL Advanced Projects Design Team's Spacecraft Instrument Cost Model: an Objective, Multivariate Approach

K. Warfield and K. Roust
Jet Propulsion Laboratory[†]

Abstract

This paper describes the development of a new satellite instrument cost model currently in use with the Jet Propulsion Laboratory's Advanced Projects Design Team. Data selection and normalization are discussed. A cost estimating relationship is developed using linear multivariate techniques in a step-wise parameter selection process with both qualitative and quantitative data. The rationale for the choice of this method over a segmented, univariate approach or a nonlinear approach is also stated. A discussion of the statistical model validation is included. A brief example and explanation of the model's use with Monte Carlo probability density functions, is covered.

Introduction

As part of the ongoing NASA effort to bring more value to the American public in this new age of reduced government spending, the Jet Propulsion Laboratory (JPL) has been undergoing an institution-wide effort to examine, and where appropriate, to re-construct its entire spacecraft development approach. One of the earliest and perhaps most successful examples is found in the arena of advanced project design. In 1995, JPL launched its Advanced Projects Design Team (also known as "Team X") as an experiment in concurrent engineering. The team consists of fourteen engineers representing all the technical disciplines necessary for the completion of a high-level mission design -- including cost. All disciplines are interlinked via a group of networked work stations running discipline-specific software applications and exchanging data through common spreadsheets. Typical tasks include trade studies, proposal evaluations, and technology road mapping. Since inception, Team X has conducted more than 200 studies most of which included at least one cost estimate. Prior to the advent of Team X, parametric cost modeling at JPL was largely relegated to the business of proposal cost validation. Only a handful of estimates was produced annually. These early parametric models were large algorithms requiring considerable subjective input and as such were deemed unsuitable for the Team X environment where cost modeling time is limited and subjective inputs are open to aggressive criticism and second guessing. This design pace pushed the parametric cost model out of the realm of institutional safeguard and into the world of product

[†] The Jet Propulsion Laboratory is operated by the California Institute of Technology for the National Aeronautics and Space Administration

design. In this re-engineering experiment, parametric cost modeling has moved from a design check to a design tool.

A new total project, total life-cycle parametric model was developed for this concurrent engineering environment. The total project model and its supporting spacecraft model are discussed in other papers presented at this conference; this paper will address the newly updated instrument cost model used for payload cost estimation by Team X.

The type of work being estimated by Team X presents several unusual challenges. NASA and JPL exist to do that which is too financially risky to attract~~ed~~ the private sector. As such, we routinely build one-of-a-kind spacecraft, often utilizing the latest technology. The data available for instrument modeling may appear meager and exceedingly diverse when compared to data from other industries or even the commercial side of the satellite business. Add to that the rather limited design definition one is likely to find in a proposal or feasibility study, and you have the Team X estimation environment. On the plus side, the typical Team X customer is a senior scientist or engineer often possessing skills in parametric estimation and as such is better positioned to understand our methods and results than someone without a working knowledge of statistics.

Aside from the usual desire to produce as high a correlation between model and data as possible, there were three other goals central to this development effort: the need to maintain a tangible connection between our estimate and the model's database, the need to evaluate uncertainty in our estimates, and the need to remain objective. We view all estimates as essentially analogy estimates. The single-point analogy method is obviously so. The ground-up, grass roots method compiles costs based on expert opinions of labor and material requirements which are themselves based on analogous experiences in that particular expert's personal history. The parametric approach is simply an analogy to the model's entire supporting database. It is for this reason -- that we consider our model estimates to be analogous to our model's database -- that we must be able to express this connection (along with our database's relevance to our customer's design) to establish the significance of our estimate to the customer's instrument. The chief method is a common scatter plot of the data and the estimate – simple, yet surprisingly effective at defusing the off hand dismissal of a parametric estimate that does not support expectations. By characterizing the uncertainty associated with an estimate, we hope to express what is known from the data and no more. There can be considerable pressure to interpret an estimate as supporting one view or another. A single point estimate can play into this politicization of the estimate by the ambiguity in its accuracy. An arbitrary, poorly supported, error figure is often attacked by review boards and rightly so as it can leave the customer with an understated understanding of the uncertainty and its associated risk. Our aim is to characterize both the model's inherent uncertainty due to residual error and sample size, and, to the extent that it is known, our customer's uncertainty with respect to the input parameters, by using Monte Carlo techniques. Our outputs are expressed as probability distributions rather than single point estimates. Lastly, objectivity is our most important aim. Aside from the speed of assessment, defensibility of estimate, calculable error, and political isolation that it brings, there is the practical reality that waging a subjective evaluation of parameters like heritage, complexity and technological development, with creditable consistency, would require in-depth instrument knowledge spanning all technologies represented within the model, and as such would be beyond the capabilities of any one person.

Data

Ultimately, the Advanced Projects Design Team Instrument Cost Model (APDTICM) rests on the specific instruments used to calibrate it. The present version of the model is based upon ninety-five instruments built generally in the 1980s and 1990s for flights in the late 1980s through the late 1990s. The selected instruments were chosen primarily based on the availability and accuracy of their cost and parametric data. For example, only instruments likely built and accounted for using full-cost accounting are included.

These missions have at least one instrument in the APDTICM database:

- ADEOS
- ASTRO-1
- Cassini
- COBE
- EOS-AM
- EOS-PM
- Galileo
- GRO
- Hubble Space Telescope
- ISTP-Polar
- ISTP-Wind
- Magellan
- Mars Global Surveyor
- Mars 98 Lander
- Mars Observer
- Mars Pathfinder
- NOAA-K
- San Marco D
- SIR-C/X-SAR
- TOPEX
- UARS
- Ulysses
- USML-1

The APDTICM instrument database initially contained nearly 40 parameters for each of over 110 instruments. As the model was built, the number of important parameters was trimmed down to six, and a few instruments were excluded from the database, typically due to age or inconsistencies in the data.

A few instruments were removed from the data set for other than age-related reasons. Two instruments (TOPEX laser reflector and LAGEOS) were removed because they were basically just mirrors, and as such, represented a wide technical departure from the body of instruments in the data. Cost data pertaining to the CUBIC instrument came from a superficial description on its web page and could not be independently confirmed. Three instruments (Cosmic Dust Analyzer and Magnetometer on Cassini and the Alpha-proton X-ray Spectrometer on Mars Pathfinder) were removed from the JPL data because the y were built and managed by foreign entities; we had concerns that these costs were not complete.

The data for these instruments were collected from many sources, including databases for previous instrument cost models, NASA databases, and local data stores. The majority of the instruments included here, some 42 data points, were also included in the 1993 Scientific Instrument Cost Model (SICM) / Multi-Variable Instrument Cost Model (MICM) database. This was our default data source. Costs for 39 instruments and some programmatic information were taken from JPL internal records. Some additional programmatic data was extracted from the 1993 NASA Cost Model (NASCOM) database and the on-line National Space Science Data Center (NSSDC). Where necessary, appropriate unit conversions were applied for the technical data. The cost data was inflated from (as appropriate) FY92, FY94, FY95, or FY96 dollars to FY98 dollars using the 16 April 1997 NASA New Start Inflation Index.

In addition to the NSSDC and MICM/SICM data, a few other data sources were used. These included some generally available data (usually non-cost data collected off of the instrument WWW home pages) and some data from NASA-only data sources (such as the REDSTAR archive run by MSFC and SAIC).

Four cost data points, related to the Mars '98 Lander mission, were taken from a UCLA proposal on the WWW. While this is not the most creditable source, the data represents only a few points, and gave no statistical evidence (such as t-ratios) that indicated that these particular data departed from the body of instruments. For the time being, we have elected to include them, but will review and possibly revise once more creditable NASA sources become available.

In developing this model, we have attempted to use quantitative data as much as possible. Quantitative parameters have two significant advantages over qualitative data. Primarily, quantitative data offers a continuum of values, which can then be analyzed alongside the continuum of cost data. Quantitative data, typically, is also clearly associated with some measurable property of the instrument (such as mass), and hence is typically not subjective.

Our qualitative data is objectively stated as a group of yes-or-no questions. The yes/no answers are then transformed into ones and zeros. These are commonly referred to as dummy variables [Draper and Smith, 1981, § 5.4] and are the usual way that qualitative data is addressed in ordinary least squares regression. Since qualitative data is not defined over a continuum, the standard, gradient dependent, "hill climbing" techniques used to find residual local extrema in the most common non-linear least squares techniques are not employable.

In keeping with the goal of retaining connectivity between the database of our model and a customer's design, the age of our data points is always a concern. Changes in technology, project implementation, agency focus...even accounting practices can affect the relevance of older data to a new instrument under study. We selected a minimum baseline launch date of 1988 or later for our model data. This avoided the use of pre-Challenger payloads that were built before the agency-wide policy impact that followed the 1986 accident, but still allowed sufficient data for our modeling.

Model Development

Model development begins with the selection of model structure. The selection of model form is partly driven by the number of data points available, number of independent variables, and the apparent behavior of the data. As this exercise was an update to an existing model, the form had already been determined, but we will review the basis of this form for completeness. The use of a multivariate model had always been taken as given; mass was well correlated to cost, but not to the point where it alone could be used for meaningful cost predictions – other parameters would need to be introduced. Segmentation of the data base into several mutually exclusive models based on, say, "instrument type" (as in the SICM model [PRC, 1990]) had been ruled out, due in part to the difficulty in defining newer instruments into single instrument types, and in part in the interest of characterizing common cost drivers on as wide a database as possible. This was accomplished by the addition of a mutually exclusive dummy variable array covering the five instrument categories (i.e., electric fields & magnetic fields, X-ray & charged particles, microwave, laser, and electro-optical) described in the Goddard MICM report [Dixon and Villone, 1990, pp. 30-31], to our multivariate model. This approach forces the model to fit all instrument categories with common cost/mass and cost/life sensitivities at the expense of a lower R^2 than may be possible with a segmented data set. While most likely not optimal, this arrangement requires less data for modeling, allows us to compare the relative importance of different independent variables, including instrument type, for predicting cost, and, as the model's R^2 is 84.2%, it still represents a reasonably good cost forecasting tool. Initially, we did try a dummy variable array using all fifteen MICM instrument families but most families did not produce convincing t-ratios, probably due to too thinly populated groupings; better results were achieved with the five instrument categories – the next higher grouping in the MICM instrument hierarchy. To acquire enough data to reconstruct the 15 original SICM categories would require going back one or two additional decades into the historical record. The advantage, if any, to segmentation did not measure up to the loss in connectivity between our data base and our customer's design that we would endure should we opt to expand the data substantially prior to the last ten years and include instruments built with older technology under outdated programmatic constraints. The choice of a linear model over a non-linear model is also a bit of a judgment call. Linear ordinary least squares restricts the model to forms that are linear, or intrinsically linear (i.e., non-linear relationships that are transformable into a linear relationships). Not all possible parametric relationships can be expressed this way. Non-linear models do not place this restriction on model form, however, non-linear models lose the clear geometric description available in the linear multivariate model and, along with it, a number of modeling measures valuable in the evaluation of predictors and model performance. Additionally, as noted earlier, non-linear analysis of qualitative data, such as "instrument type", is particularly problematic due to the discontinuous nature of the parameter and the local minimum solving techniques found in standard non-linear methods [Björck, 1996, § 9.1.2]. We elected to presume a linear fit until a time when the statistical evidence compelled us to do otherwise.

The model was generated using standard least squares fits to linear, log-linear, and log-log forms on a trial-and-error approach. The data were randomly divided into two groups: one group of 65 points was used for model regression, and the other group containing the remaining 30 points was used for model validation. The statistics package used was MINITAB®. MINITAB® allows the predictors to be loaded in and regressed in bulk against the dependent variable by systematically going through all predictor combinations and printing out those with the highest Pearson coefficients for each number of predictors. This substantially reduces the workload and permits the evaluation of hundreds of predictor combinations as possible models. Once likely predictors have been identified, a detailed regression is run to generate model coefficients, coefficient t-ratios, cross-correlation measurements, and an overall R^2 value. Different runs are compared with the objectives of maximizing R^2, minimizing cross-correlation and providing predictors with creditably high t-ratios. This last part is somewhat arbitrary, but about two standard deviations were used as the minimum acceptance criteria. The final model equation and associated statistics (after transformation into linear form) are shown below in Figure 1. Cost is in 1998 millions of dollars, mass is in kilograms, and design life is in years. All four dummy variables are valued as 1 for "yes" and 0 for "no".

$$\text{Cost} = 2.5883*(\text{Mass}^{0.567})*(\text{Life}^{0.439})*(0.3133^{\text{EF\&MF}})*(0.4112^{\text{XRAY}}) *(0.3607^{\text{COPY}})*(0.6965^{\text{UNIV}})$$

Variable	Model Data			Variable Coefficient		
	Mean	St. Dev.	Range	Coef.	St. Dev.	t-ratio
ln(Cost)	2.811	1.351	-0.5 to 5.5	--	--	--
Constant	--	--	--	0.9511	0.2255	4.22
ln(Mass)	3.567	1.925	0.3 to 9.4	0.5673	0.0416	13.65
ln(Life)	0.994	1.064	-2.3 to 2.1	0.4389	0.0740	5.93
EF&MF	0.123	--	--	-1.1607	0.2173	-5.34
X-Ray	0.308	--	--	-0.8887	0.1573	-5.65
Copy	0.092	--	--	-1.0196	0.2427	-4.20
Univ.	0.246	--	--	-0.3617	0.1580	-2.29

ANOVA:

SOURCE	DF	SS	MS	F
Regression	6	100.142	16.690	57.85
Error	58	16.734	0.289	
Total	64	116.876		

$R^2 = 84.2\%$

Figure 1: Model Statistics

The independent variables selected as significant cost drivers were mass, mission life, and four dummy variables: electric & magnetic field instruments, X-ray & charged particle instruments, university built instruments, and copied instruments. Mass has long been identified as an instrument cost driver -- that it is the dominant component in this model is not unexpected (see also Figure 2). The relation to mission life has been suspected for some time, but is often handled with the surrogate of mission type, as is the case with the 1990 MICM model [Dixon and Villone, 1990, p 8] and the earlier version of this model. We view the connection of cost to mission life rather than cost to mission type as the more direct and meaningful, and as such, it represents a major improvement over our earlier work. The two instrument category dummy variables are the same as used in the original version of this model. The lower cost per kilogram associated with either of these two instrument categories is most likely due to less sophisticated technology than found on other instruments in general. Since the costs attributable to these parameters is the component exclusive of mass, they are probably the closest objective representation of complexity we have yet uncovered. That copies of previously built instruments cost less than new designs should surprise no one. This is, of course, due to the lack of design and other non-recurring costs in the payload. This parameter was missing from the original model, so its inclusion will allow better estimation of second unit costs than was previously possible. And lastly, we ascribe the lower university-built costs to the lower labor rates typically found in educational centers where work forces are often augmented with students, and post doctoral staff members.

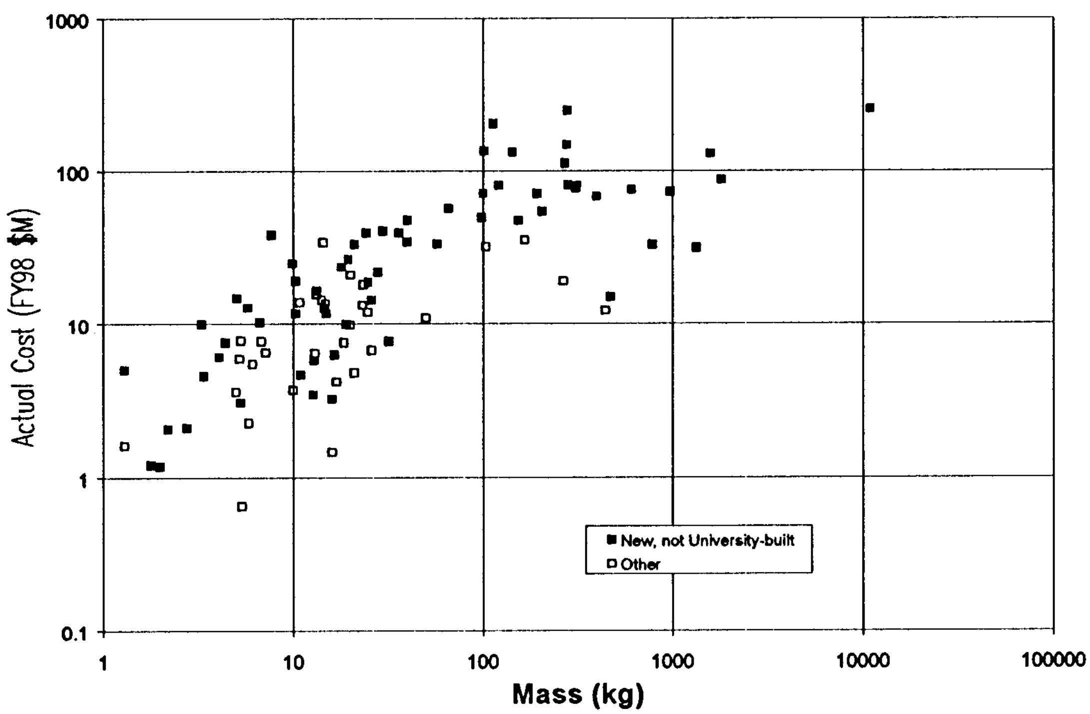

Figure 2: Mass vs. Actual Cost

The total cost includes design, development, test and evaluation costs plus the first flight unit cost. The total cost does not include contractor fees or instrument to spacecraft integration costs. The total mass includes the mass of all parts of the instrument, including such things as antennae and telescope optics as are required for operation. The design life is the length of time that the instrument was designed to operate. For most instruments, it should be roughly the period of time between launch and the end of life for the spacecraft.

An assumption was made that typical Space Shuttle instruments (which fly for approximately two weeks) are designed to last 0.1 years. This assumption is likely true even for instruments designed to fly on several Shuttle missions, as simple components can be replaced between flights. The instrument categories are identical to the MICM instrument categories [Dixon and Villone, 1990, p. 30]:

- Electric Field and Magnetic Field Instruments -- This instrument group detects/measures magnetic and electric fields within the magnetosphere, outside the magnetosphere, between the geo-magnetic and interplanetary fields -- plus fluctuations in these fields. Magnetic field instruments have search coil, fluxgate, and atomic nuclei sensors. Electric field instruments have antennas and/or paired-probe sensors. Included in each type of instrument are associated electronics with detection/analyzer and data processing capability.

- Charged Particle and X-Ray Instruments – This instrument group detects/measures the properties of X-Rays/Gamma rays and/or charged particles in the plasmas that have been trapped in the magnetospheric region of earth or another planet. The particle measurements include type, number, mass, and kinetic/thermal energy. There are three instrument classes: (1) Instruments that detect X-rays/Gamma rays and/or cosmic ray particles by their interaction with matter; (2) Instruments that determine particle mass with an electromagnetic analyzer; and (3) Instruments that detect/measure kinetic/thermal energy with an electrostatic analyzer.

Instruments that are exact copies include only instruments that are built identically to previously built instruments. Very few instruments should be exact copies, and the scientists and engineers responsible for those that are exact copies should know it. For example, the Mars Observer instruments that were rebuilt and flown on Mars Global Surveyor and Mars 98 Surveyor are exact copies. If any significant changes are made during rebuilding or retrofitting an instrument (like changing to electronics that are more modern), then this likely is not an exact copy. The exact copy question is meant to capture instruments that have no new engineering work. The "instruments built by an educational institution" parameter is supposed to capture some of the cost savings which educational institutions seem to gain from using student labor. If we divide the world into three types of entities: industry, Federally Funded Research and Development Centers (e.g., JPL, LANL), and educational institutes, only instruments built entirely by the last group should answer 'yes' to this question.

Notable parameters tested but not selected for use in this model were power, data rate, launch year, and lead NASA center. Data rate, launch year and lead center were dropped because low t-ratios implied that they had little effect on cost. This is reassuring in the case of lead NASA center as a particular stand out could be interpreted as an indication of irregular accounting methods. Power data did have a significant t-ratio but only raised R^2 by 1% and severely lowered the mass t-ratio. This is a sign of significant correlation between mass and power (see also Figure 3). Since the improvement was marginal and bought with the removal of several data points lacking power information we elected to omit power from our model.

The mass parameter data ranged from 1.3 kg to 1817.7 kg (and one instrument -- SIR C -- at 11 metric tons). Even after transformation, the SIR C mass appears as an outlier, however, in all other characteristics SIR C is comparable to the rest of the data base. We have decided to include it, as it fits the model and represents a rare historical example of an extremely heavy payload, but model estimates far above 1800 kg are more anecdotal in nature and should be handled in another manner if possible. The design life data ranged from 0.1 years to 8 years. It is important to note these ranges as they define the limits of

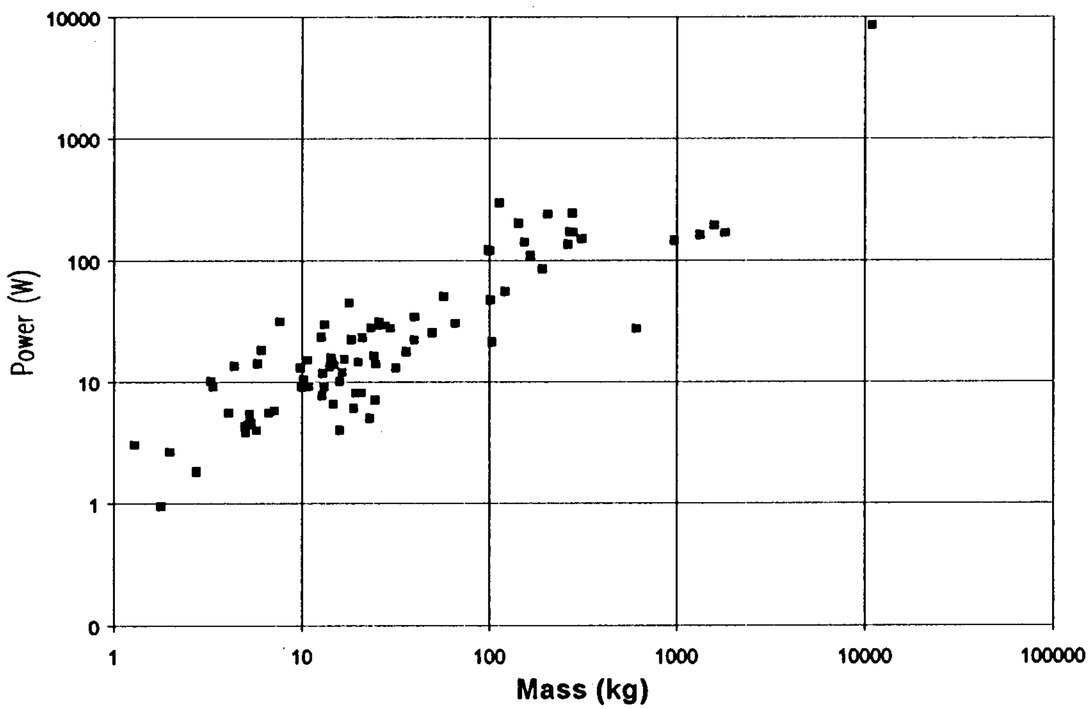

Figure 3: Mass vs. Power

the model. Input data outside these ranges can produce dubious results and are indicative of a lack of relevance of our database to that particular design.

Residual Analysis

Next, the residual for this equation was examined graphically. It should be noted that since our model is multiplicative and is transformed into linear by taking the logarithm of the model, we are not evaluating the multiplicative error ε, but rather the natural logarithm of the error, $\ln(\varepsilon)$. Plots of the residual versus each independent variable were run to verify uniformity over the entire model range (see Figure 4, Figure 5, and Figure 6). An

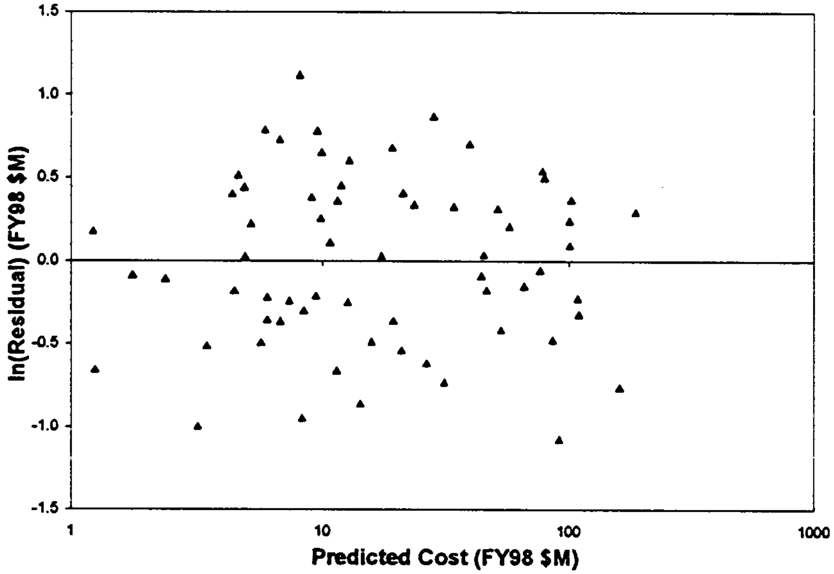

Figure 4: Predicted Cost vs. Residual Cost

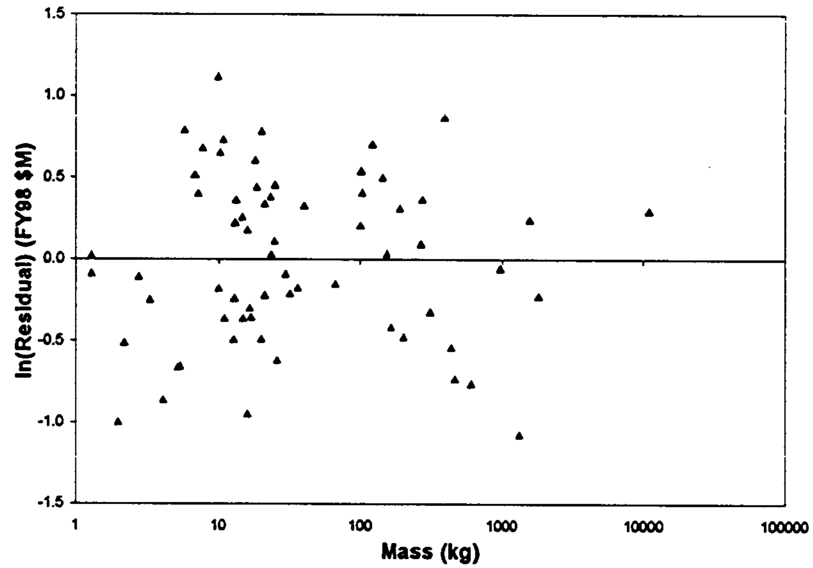

Figure 5: Mass vs. Residual Cost

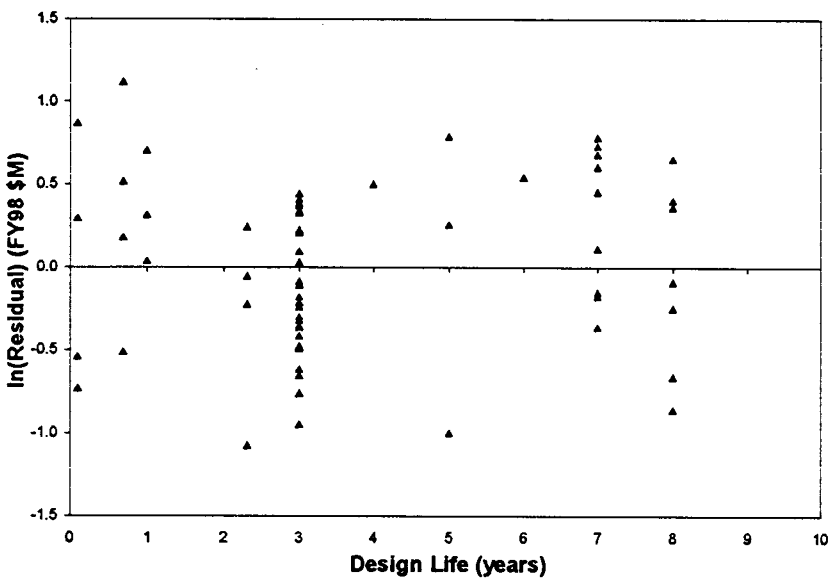

Figure 6: Design Life vs. Residual Cost

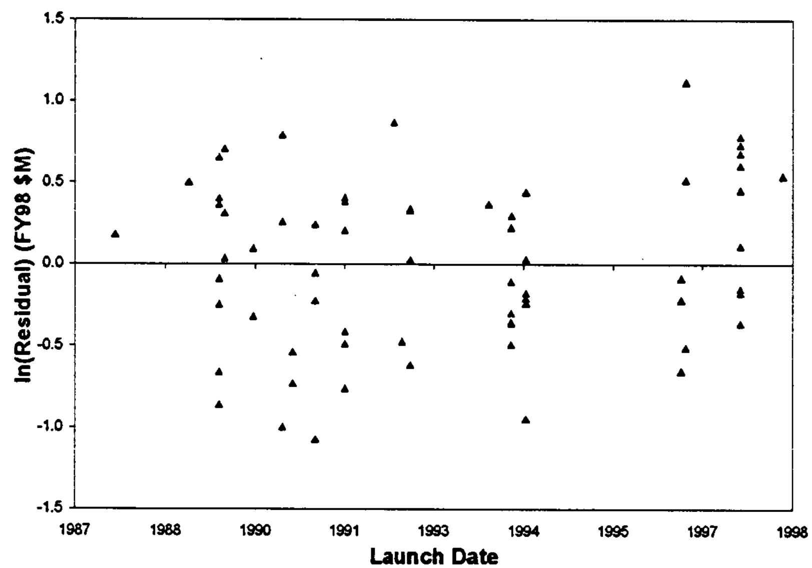

Figure 7: Launch Date vs. Residual Cost

additional plot of residual verses launch date was also run (see Figure 7) to evaluate if any uncaptured, time based cost components exist: our conclusion is there is no discernible cost trend over our ten year baseline. A histogram (see Figure 8) of the residual was also generated as a crude gauge of the error distribution (a fundamental assumption in multivariate modeling is that the residual is normally distributed about zero). A more rigorous confirmation was carried out by first demonstrating that residual mean was within sampling error of zero using

$$z = \frac{(\mu_r - \mu)}{(\sigma/\sqrt{n})}$$

where n is the sample size, μ_r is the residual mean, μ is the normal population mean, and σ is the normal population variance. Since our model's residual is standardized, the comparable normal distribution has mean equal to zero and a variance of 1. With a residual mean of -0.002 and a sample size of 65 points, the model's standardized residual has a z number of -0.016 with respect to the standard error of the mean of the normal distribution; essentially the model mean is easily within sampling error to the hypothesized normal mean of zero. Two other tests were performed to check the residual normal distribution requirement: the first was a comparison of the residual and normal population variances using the F test; the second was a comparison of the two population distributions using the Chi-square test. The first test is simply taking the ratio of the two variances and comparing it to a critical value from an F distribution for a chosen confidence level [McClave and Benson, 1988, § 9.3]. In all our tests, we use a 95% confidence level. If the ratio is above the critical value, we accept the null hypothesis that the two samples represent different populations. Conversely, if the ratio is below the critical value, we must reject the hypothesis. With an F test statistic of 1.018, the variance ratio of the standardized residual and normal distribution, is under the critical value of 1.53. There is no evidence to suggest that the distribution of the residual is anything but normal. This test presumes that both sample populations are normally distributed. In the case of the Chi-square test, however, this restriction is not required. The Chi-square test compares the two distributions against a critical value selected from a Chi-square distribution, and as in the F test, values above the critical value indicate acceptance of the null hypothesis and values below indicate rejection. Distributions are characterized by segmenting into a group of cells. The number of data points within each cell are counted and compared to the number of data points within the corresponding cell in the other sample distribution. Cells must be selected so that no cell has less than five counts. The test statistic is [McClave and Benson, 1988, § 17.1]:

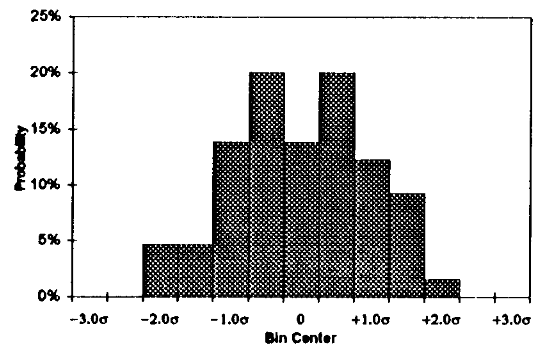

Fig. 8: Histogram of Model Residual

$$X^2 = \sum \frac{[n_i - E(n_i)]^2}{E(n_i)}$$

where i is a particular cell in the distributions, $E(n_i)$, the expected number of outcomes, is the number of data points from the normal distribution, and n_i, the number of outcomes, is the number data points from the residual distribution. With a Chi-square statistic of 2.246 and a critical value of 12.592 we again conclude that there is no evidence that the residual is other than a normal distribution.

Model Validation

The thirty data points not used in model generation, were run through the model to create a population of predicted values. Validation was carried out by comparing predicted and actual cost populations with F and Chi-square tests similar to those used to compare residual and normal populations earlier in this paper. The validation F test statistic was 1.015 compared to a critical value of 1.84, and the validation Chi-square statistic was 5.244 compared to a critical value of 7.814. With a 95% confidence level we conclude that there is no evidence that predicted and actual populations differ.

Model Output

Once the model had been verified, we needed it in a form that could be easily used. To that end, we built a workbook in Microsoft Excel® that would be able to accept inputs from the user and return the cost estimate (for the mean cost), along with a confidence interval and a prediction interval. Confidence and prediction intervals were calculated using the following linear least squares equations [Johnson and Wichern, 1992, pp. 305-306]:

Confidence Interval: $\mathbf{z_0}'\beta \pm t_m(\alpha/2)\sqrt{s^2(\mathbf{z_0}'(\mathbf{Z'Z})^{-1}\mathbf{z_0})}$

Prediction Interval: $\mathbf{z_0}'\beta \pm t_m(\alpha/2)\sqrt{s^2(1 + \mathbf{z_0}'(\mathbf{Z'Z})^{-1}\mathbf{z_0})}$

where $\mathbf{z_0} = [1, z_{01}, \ldots, z_{0r}]$ is a vector of input values for the r model independent variables, β is the model coefficient vector, $\mathbf{Z}$ is the design matrix, s is the standard error of the estimate, and $t_m(\alpha/2)$ is the upper $100(\alpha/2)$th percentile of a t-distribution with m degrees of freedom.

We then attached graphics software (for scatter plots of the data and estimate), which allows us to get a better feel for the actual distribution of the data (see Figure 2 for an example). Finally, in keeping with our goal of being able to express the error associated with our estimates in a quantitative and understandable method, we installed a Monte Carlo generation add-in to the Excel-based model to create histogram probability density functions (PDFs) for each estimate. A brief description of the process is given by Sobol' [1994].

In our Monte Carlo simulation, uncertainty in the design is captured with PDFs for the input parameters, and uncertainty in the model due to the standard error of the estimate and uncertainty in the model coefficients is represented as a normal distribution with mean equal to the model prediction and standard deviation equal to the square root of the variance of the forecast error. These uncertainties are convolved to produce an output PDF for the cost estimate. An example using triangular distributions for mass and mission life is given in Figure 9 and Figure 10. In particular, we used a triangular mass distribution with a minimum value of 45, a most likely value of 50, and a maximum value of 52 kg. The design life distribution was triangular, with a minimum of 1.5 years, a most likely value of 2, and a maximum value of 3 years.

Conclusions

Through a multivariate approach, we were able to create an instrument cost model with reasonably high coefficient of determination, calculable uncertainty, and requiring only minimal design knowledge. The relative strengths of cost drivers are now readily identifiable through model coefficients and t-ratios. Objectivity, recognized statistical techniques, and graphical tools enable us to convey the connection from our database, model, and estimate, to our customer's design. The final product is an instrument cost model valid over a wide range of instruments, masses and mission durations, and capable of the rapid, objective estimates needed in the concurrent engineering setting.

Acknowledgments

This work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, under contract with NASA. The authors wish to thank George Fox for his Monte Carlo simulation tool, and L. Rosenberg, G. Madrid, H. Habib-agahi, and J. Hihn for their comments and suggestions.

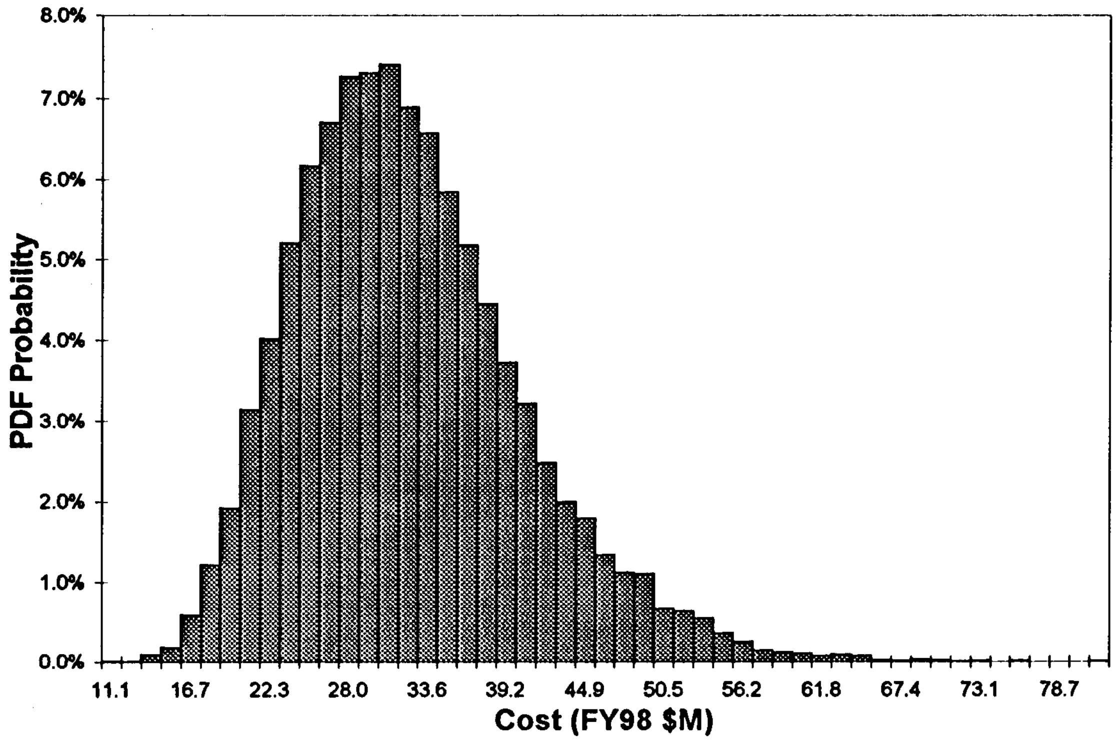

Figure 9: Example Probability Density Function for Instrument Cost

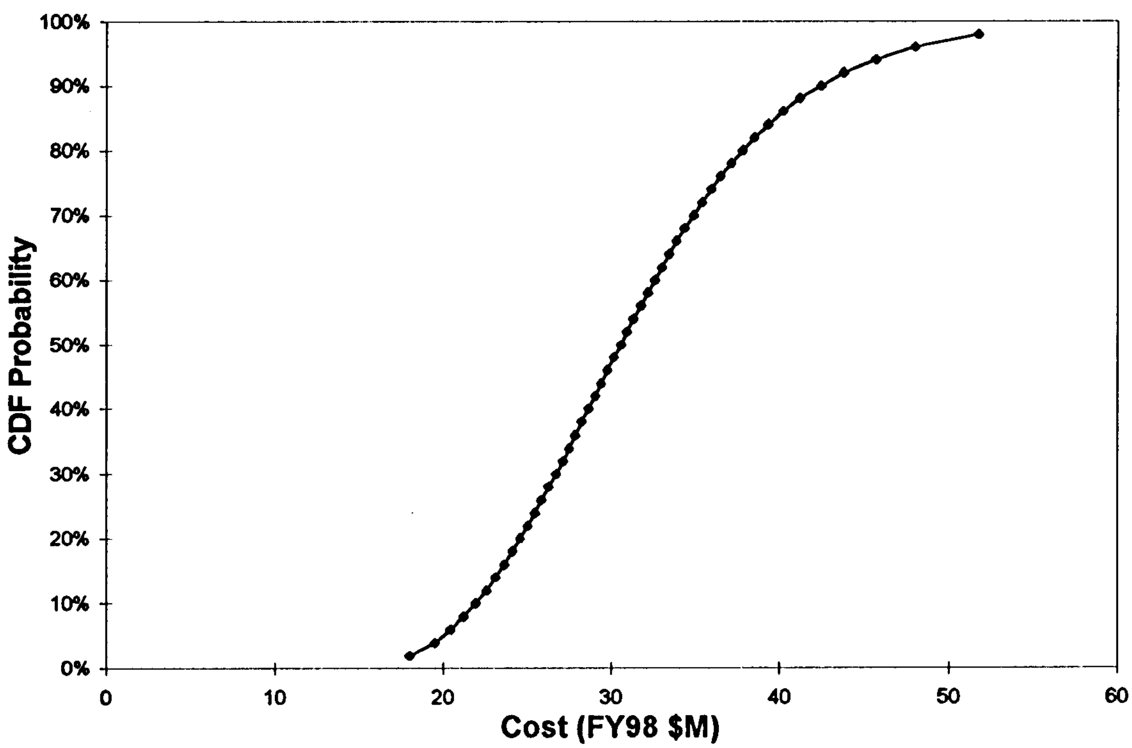

Figure 10: Example Cumulative Probability Density Function for Instrument Cost

References

Björck, A., *Numerical Methods for Least Squares Problems,* Philadelphia: SIAM, 1996.

Dixon, B. and P. Villone, Goddard Multi-Variable Instrument Cost Model (MICM), RAO Note #90-1, May 1990.

Draper, N. R. and H. Smith, *Applied Regression Analysis,* New York: John Wiley, 1981.

Johnson, R. A. and D. W. Wichern, *Applied Multivariate Statistical Analysis,* Englewood Cliffs, New Jersey: Prentice Hall, 1992.

McClave, J. T. and P. G. Benson, *Statistics for Business and Economics,* San Francisco: Dellen Publishing, 1988.

PRC Systems Services, Scientific Instrument Cost Model (SICM), Technical Brief 125, PRC D-2327-H, February 1990.

Sobol', I. M., *A Primer for the Monte Carlo Method,* Boca Raton, FL: CRC Press, 1994.

Author Biographies

Keith Warfield is a Senior Project Cost Analyst at the Jet Propulsion Laboratory working with the Advanced Projects Design Team. He holds a B.Sc. in Engineering and Applied Science from the California Institute of Technology. Prior to his work as a cost analyst, Mr. Warfield worked for more than a decade in engineering roles on several science instruments including the Satellite Test of the Equivalence Principle, the Mars Observer Pressure Modulator Infrared Radiometer, the TOPEX Microwave Radiometer, the UARS Microwave Limb Sounder, and particle detectors for the Stanford Linear Accelerator.

Kevin Roust is a Project Cost Analyst at the Jet Propulsion Laboratory. He received a B.Sc. in Economics from the California Institute of Technology. Mr. Roust has been involved in cost tool development for the Advanced Projects Design Team, and other computer simulation work for JPL's re-engineering effort.